

THE COMPUTER SCIENCE RESEARCH NETWORK CSNET: A HISTORY AND STATUS REPORT

DOUGLAS COMER Purdue University

1. THE HISTORY OF CSNET

Douglas Comer's current research interests include operating systems, and computer communication networks.

Author's Present Address:
Douglas Comer, Computer Science Department, Mathematics Building, Purdue University, West Lafayette, IN 47907.

The CSNET story began in Madison, Wisconsin on a spring morning in May 1979. Lawrence Landweber, then chairman of the University of Wisconsin's Computer Science Department, had invited a group of colleagues from other universities as well as representatives from the Defense Advanced Research Project Agency (DARPA) and the National Science Foundation (NSF) to discuss "the feasibility of establishing a Computer Science Department research computer network." For two days, the 13 participants, representing six universities and two potential funding agencies, heard presentations and exchanged views. All were excited by the prospect of a computer science network which would connect their organizations. Yet each held a different view of what the network should provide, how it would operate, and how it could be funded. As they sat listening, past experiences came to mind.

Robert Kahn, director of the IPTO office of DARPA, had contributed to the early research and development of ARPANET, one of the oldest long-haul, packet-switched networks; he administered a research program that kept ARPANET at the state of the art. The group knew that ARPANET provided services like mail and file transfer, sites with ARPANET access engaged in research for the Department of Defense, and ARPANET connected military sites, government sites, and a handful of universities. They also knew that because communication over the network was convenient, researchers with ARPANET access tended to interact among themselves and ignore those without access.

Peter Denning of Purdue University appeared at the Madison meeting, but not as an observer. As vice president of the ACM and head of a large Computer Science Department that did not have ARPANET access, Denning was acutely aware that the top computer science departments, such as those at

ABSTRACT: In 1981, the National Science Foundation started a five-year project totaling nearly \$5 million to construct a computer science research network, CSNET, connecting all groups engaged in computer science research. For an NSF division with an annual budget of \$25 million, the award represents an unusual commitment to a single project; only a handful of such large awards have been made. What is CSNET? Why is it receiving such attention? How will it benefit the community? When will it be completed? Who are the architects and implementors?

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission. © 1983 ACM
0001-0762/83/1000-0747 \$5.00

The Massachusetts Institute of Technology, Carnegie-Mellon University, and Stanford University, used ARPANET extensively. Like many others at the meeting, he sensed a growing split in the community between those with ARPANET access and those without it. He believed that network access would be critical to all computer science departments in attracting and retaining faculty and students.

Purdue was typical of many non-ARPANET schools: They had joined with Bell Laboratories and several other universities to exchange messages and small files over telephone lines. In the beginning, the research computer at Bell Labs polled all sites on this informal network once a day over dial-up phone lines, retrieving all outbound messages, and depositing messages that it had queued waiting for that site. The software, which ran under the UNIX¹ operating system, was known as uucp (UNIX to UNIX copy program). Phone calls were expensive, and uucp traffic had built up to the point where it began interfering with other uses of the computer. In the spring of 1979, Bell Labs announced that it would no longer transship mail between non-Labs sites; when the network service ended, it left several universities disconnected.

Anthony Hearn, then head of the Department of Computer Science at the University of Utah, was among the few attendees who had used ARPANET first-hand. He had organized a group of researchers who shared his interest in symbolic computation into a group that exchanged information and programs over the network. He was interested in expanding SYMBOLNET to include researchers at universities without ARPANET access. Realizing that telephone connections were slow and expensive, Hearn had started a project to use public packet-switched networks to reach beyond ARPANET.

Landweber held another view of networks. Although his primary research had been in theoretical computer science, he realized that network communication could significantly improve his research environment by allowing him to better communicate with fellow researchers. As a department chairman, he too believed that without network access, researchers in his department would lose ground to those in other departments.

Two years earlier, after discussions with his colleagues Richard DeMillo and Richard Lipton, Landweber had organized an electronic mail facility for theoreticians; it was called THEORYNET. Funded by NSF, THEORYNET provided members a mailbox on a central computer at the University of Wisconsin which users accessed from terminals over dial-up phone lines or through the Telenet public packet-switched network. Although THEORYNET had been successful in attracting users (over 100 researchers working in the area of theoretical computer science had sent or received messages), it was limited to message transfer. Landweber wanted to extend THEORYNET to include file transfer, remote login, and faster message delivery. Such extensions required radical change: Instead of connecting terminals to a computer, the new network had to connect computers to other computers. It would require every site to agree on the means of connection and the communications software. Landweber hoped that the group could agree on the technology quickly and move on to implementation.

Kent Curtis, head of NSF's Division of Mathematical and Computer Sciences, was not surprised to receive an invitation from Landweber to attend the Madison meeting. NSF had funded Landweber's THEORYNET project; the two talked frequently.

¹UNIX is a trademark of Bell Laboratories.

Curtis arrived in Madison with much more than THEORYNET on his mind. He knew that many computer scientists in universities depended on NSF for their research funding, and that the split between those with ARPANET access and those without was widening. He also knew that industries had been luring faculty and students away from basic research at universities with high salaries and superior facilities. The migration away from universities had come at a time when the national demand for computer specialists to train students was at an all-time high and still increasing. Eventually, the declining supply of trained students and increased national demand could trigger a crisis.

Working from reports that blamed part of the problem on the inadequate and outdated facilities available in most universities, and from the advice of leaders in the field [1], Curtis was planning new NSF programs that would help schools improve their facilities. Some of that advice came from NSF's Computer Science and Engineering Advisory Panel in a resolution adopted four years earlier [2]:

We recommend that NSF provide to qualified computing researchers easy access to an international computer network. This access would create a frontier environment which would offer enhanced communication, collaboration, and the sharing of resources among geographically separated or isolated researchers.

1.1 The Proposal

After listening to the discussion, Curtis agreed that a network could make a significant improvement in university facilities. He suggested that the group submit a proposal to NSF for evaluation through peer review. He offered a few guidelines: (1) NSF wanted to help many schools; (2) the network should provide a variety of services with cost proportional to use; (3) NSF funds were limited—at most, \$5 million would be available for the entire project; (4) NSF would not pay costs indefinitely, so the network should be self-sustaining after five years; and (5) NSF sought community participation in the design and construction of the network.

During the summer of 1978, Landweber, with the help of Denning, Hearn, and Robert Ritchie (chairman, University of Washington), sketched out a proposal. They quickly realized that \$5 million was not much money for what the group envisioned. Besides the cost of equipment and software, the money had to cover the cost of managing the software development, the cost of establishing an organization to collect and distribute bills, the cost of communications between sites involved in the development, and the cost of publishing and disseminating information and software. Clearly, the new network would have to rely on existing communication services; the funds would not cover equipment cost for a direct copy of the ARPANET.

Everyone was attracted to the idea of using public packet-switched networks (e.g., GTE Telenet) for machine interconnections. Such services were available almost everywhere (including Europe), and they advertised higher speed connections than dial-up telephone. In addition, low-speed connections could be installed initially and upgraded to higher speed should the site want to spend more. So, the proposal was based on the idea of a net-connecting computer science department machines over public packet-switched networks. It was called CSNET.

The group sent copies of the draft out to the community. Reactions were mixed. Some said the proposed network would provide a quantum improvement in their environments; others questioned the rationale behind the plan. How would the protocol software be implemented? Who would

manage the development? What type of organization would be established to run the network? Why not use ARPANET instead? In December 1979, Landweber revised the draft to address these questions and submitted it to NSF, hoping for funding by the following summer. NSF sent the proposal to many reviewers in academia and industry.

The reviews came in April 1980 and were skeptical. It was agreed that the need was great, but there were reservations about the project as proposed. Would the proposed distributed management scheme work? The principal investigators were not engaged in research on networks; were they qualified for such a large project? The technology already existed; should NSF support a project with a large development component? How would researchers on CSNET communicate with those on ARPANET? NSF felt it could not commit funds to a major project while such doubts remained. NSF also felt it could not completely reject a proposal that had such strong support from some reviewers. As a compromise, NSF offered to fund further study of the problem. If the study went as expected, it would result in a revised proposal that would be acceptable to all.

1.2 The Study

At first, the news depressed those who had invested considerable time and effort in the proposal. The reviewers' comments showed that the proposal lacked a concrete plan. In their concern for non-ARPANET schools, the group had ignored ARPANET, but as the reviewers pointed out, communication with ARPA sites would be important. The proposal also excluded researchers who could not afford a connection to a public packet-switched network (at least \$1000 per month), and researchers who could not persuade their computer centers to run the CSNET software.

NSF wanted the group to understand NSF procedures and budget limitations. They appointed Charles W. Kern, a chemist who had used computers in his research, to represent NSF in the CSNET discussions. Although Kern knew little about the technical details of computer networks, his work at NSF helped him guide the group toward a better proposal. He describes his task as that of a protagonist—one who worked to assess and implement the wishes of the computer science community. Most important, he encouraged the group to consider all parts of the community including small schools, individual researchers, and industrial research labs.

At DARPA, Kahn learned that one of NSF's major concerns was how CSNET sites would communicate with ARPANET sites. Concerned about the future of computer science departments and eager to extend the cooperation between DARPA and NSF, Kahn appointed one of his scientists, Vinton Cerf, to work with the CSNET group to find a way to connect the two networks.

Encouraged by these developments, Landweber decided to regroup and revise the proposal. In a flurry of activity during May 1980, he requested NSF funding for the study, and began to rally the community around the idea of CSNET. He organized a meeting at the University of California-Berkeley attended by a larger group representing more of the computer science community.

The Berkeley meeting was important in several ways. Scientists like Bruce Arden, Alan Batson, Fernando Corbató, Jerome Feldman, Richard Fateman, and Nico Habermann attended. For many, it was their first introduction to CSNET. With the new faces came new perspectives and new ideas. DARPA announced their support for the CSNET project, and Cerf volunteered to work on an architectural plan that would

interconnect CSNET and the ARPANET. Other groups were formed to study ways to organize the business of running CSNET, and to devise an implementation strategy.

Ritchie volunteered to send out a questionnaire asking departments what services they needed and how much they were willing to spend. The results were important because they helped define the needs and the constraints. The survey suggested that schools wanted a wide range of services including message and file transfer, but some schools could not afford high-speed network connections. So the planners began looking for a scheme that would encompass more of the community by making costs depend on the speed, volume, and quality of service.

David Farber of the University of Delaware offered a solution to the problem of providing low-cost mail service for sites that could not afford direct connections to public networks. He had been running a mail relay service for the Army that used dial-up telephone connections to retrieve and deposit messages. The relay sent mail destined for ARPANET sites directly over the network, but queued mail destined for those sites contacted by telephone until they were polled. Each site could determine how frequently it was polled and pay the resulting phone cost. Sites that had dial-out equipment called the relay themselves whenever they chose. The software, called MMDF, ran under the UNIX operating system [3]; Farber offered it to CSNET without cost.

MMDF was attractive. Because the software was already written and debugged, it could be sent out immediately, providing early mail service for everyone. In the long term, it would shuttle mail to sites that could not afford packet-switched network connections. Meanwhile, vendors offering public packet-switched network services were contacted to obtain more details about availability, cost, and potential discounts. GTE Telenet, one of the largest, was chosen as the tentative vendor.

1.3 The Revised Plan

A revised plan containing three major components grew out of the work of the Berkeley meeting. The components were: general network services, protocol development, and PhoneNet.

Services. A services component was needed to define the services that the network would supply and to establish the software that would provide these services. The group envisioned several network-wide services like "directory assistance" that would help users locate one another. In addition, the services component included the work of building organizations to operate and to manage the network once it became self-sufficient.

Protocols. The protocols component was needed to define the interactions between CSNET member's machines and the networks like Telenet and ARPANET that connected them. The protocols component would study what hardware could be used to connect member's machines to the public packet-switched networks, and would establish the necessary software.

PhoneNet. The telephone relay component would provide low-cost mail service to sites that could not afford direct high-speed network connections. The term PhoneNet came into use for the set of CSNET sites constituting this imaginary "network."

As discussions continued, the components became more focused.

Later that summer, Cerf laid out his plan to provide for interconnection between ARPANET and the other CSNET

component networks. CSNET would be a logical network composed of several physical networks. Software would be built to hide the differences, making all services available to all users. Users on ARPANET would access CSNET services through ARPANET. As suggested in the original proposal, an equivalent set of services would be available to CSNET Telenet users. A computer called the Value-Added Network (VAN) gateway would connect the two into what is known as an internet. Already a leader in the area of internet connections, DARPA was working on a new set of communications protocols called TCP/IP that CSNET could use to automatically route information across the gateway onto the appropriate network. Furthermore, DARPA was willing to build the VAN gateway and to give CSNET the TCP/IP software. The planning group accepted Cerf's proposal [4].

By August 1980, when the planning group met again, reports were available from the organizational, implementation, and architecture subgroups. They had identified several goals: the cost should be proportional to the type and volume of service; all researchers should have access to CSNET; CSNET should eventually be self-sufficient; and development should cost less than \$5 million and take less than 5 years. They had a firm technical plan, including support from DARPA.

With this, Landweber and a few others began to revise the proposal. They chose the TOPS20 and the UNIX operating systems for the initial implementation (based on the number of sites that could be served). They divided the project into specific tasks and designated a manager. They worked with Kern to gather and document community support, and with DARPA to decide how duties and responsibilities would be divided. They picked five sites: the University of Wisconsin, Purdue University, the University of Delaware, The Rand Corporation, and the University of Utah, to do the initial development. They found faculty members with systems experience to head the technical projects. Marvin Solomon and Landweber, both from Wisconsin, would head the CSNET nameserver and service machine projects, respectively. Douglas Comer of Purdue and Lee Hollaar of Utah would head the communications protocol work cooperatively.

In October 1980, the group submitted its revised proposal to NSF for evaluation, and waited again. Landweber, Denning, and others began to solicit support from key people, taking time to discuss the need for CSNET, and looking for ways to improve the project. Meanwhile, Kern assumed responsibility for soliciting reviews and the other details necessary to move the proposal through NSF.

1.4 Before the Board

In late December 1980, with a fresh set of reviews in hand, NSF told Landweber that the evaluation had gone well. However, one major hurdle remained. Projects requiring substantial funds had to be presented to the National Science Board for approval. NSF told Landweber they would submit the proposal to the Board at its winter meeting in January. If the project was not approved in January, the Board would not meet again until late spring. It had been 18 months since Landweber's first meeting in Madison; another six-month delay seemed intolerable.

The main question raised by the second set of reviewers was whether the proposed investigators had enough experience to manage a large, multi-site project. Such a project was bound to uncover interesting research problems that could easily distract a professor, they argued. They also asked how a manager at one university could keep a group at another university on schedule. As Kern prepared to present the project to the National Science Board, Kent Curtis identified the

management problem as a weak spot. He was convinced that the Board would not ignore the reviewers' objections. To guarantee adequate management, he reasoned, NSF would have to take an active role. It would be an unusual step for NSF, but there were no alternatives. Curtis appointed Kern to be a full-time project manager and oversee all the components of the project.

With only a few weeks before the meeting of the Science Board, Kern asked the group for more details. Could they be more specific about the responsibilities assigned to each site? Could they develop five-year workplans? Confident that the proposal would be approved, Landweber called groups together to develop workplans and schedules. For some, it was the first time they met.

All went well. Kern received the workplans, and presented the project to the Science Board in mid-January 1981. They approved with one stipulation—that NSF extricate itself from management by January 1983. This implied that a major part of the activity would be focused on setting up the organization to collect and disburse funds, and that the project must be sufficiently advanced after two years that users would be willing to begin paying dues and fees.

1.5 The Contracts

After the project had been approved, each site had to file and sign a contract stating its responsibilities. The University of Delaware took responsibility for PhoneNet. The University of Wisconsin took responsibility for establishing a CSNET "service" machine that would provide services like simplified network mail addressing and directory look-up (nameserver), and took responsibility for acquiring equipment and forming the CSNET organizations that would establish policies and manage the network. Purdue University and the University of Utah took responsibility for specifying and implementing communications protocol software to link DARPA's TCP/IP to Telenet's protocol, X.25.

Minor problems delayed the signing. While contracts were being drafted, lawyers debated who had the right to publish reports and how to specify the sites' workplans. Hearn moved from Utah to the Rand Corporation during the process. The management committee decided to concentrate protocol efforts on the VAX UNIX system, so the Utah work on TOPS20 was dropped; Purdue took over. Meanwhile, Hearn convinced the Rand Corporation to become involved in the telephone mail relay effort, and Rand became the west-coast counterpart of the Delaware relay.

During this period, the project was defined and documented in detail. The NSF contracts would provide seed money to build an initial version of the CSNET software, establish an initial set of services, and form organizations to manage and control the network.

Kern expedited the paperwork within NSF. He formed technical, policy, and organizational support groups with representation from the project staff and the computer science community to help review project progress and to guide in the eventual formation of an organization to run the network. He also defined a management committee and appointed people to it: Landweber (Chairman of the Policy Support Group), Farber (Chairman of the Technical Support Group), Hearn (Chairman of the Organizational Support Group), Denning (At-Large Member), and Kern (Project Manager).²

By June 1981, when the dust had settled, contracts were signed, and work began in earnest.

²Richard Edmiston was added to represent the CIC after that organization was formed (see section "Phase Three").

GEOGRAPHIC MAP, JUNE 1, 1983



2. THE STATUS OF CSNET

CSNET development proceeded on several fronts simultaneously; milestones were scheduled in roughly three phases. Phase One established mail relay services and a public host. Phase Two will implement software to connect hosts across Telenet and establish a nameserver. And Phase Three will complete the formation of organizations to control and manage the network, making it self-sufficient. According to the schedule, Phase One was to be completed within the first year. Phase Two is scheduled to be completed at the end of three years. And Phase Three is scheduled to be completed at the end of five years. The development contracts reflect this schedule; those concerned with Phases One and Two end after three years, while others end after five years.

2.1 Phase One—Public Host and PhoneNet Relays

By July 1982, Phase One was complete. The public host, temporarily located at the University of Wisconsin, provided dial-up access to CSNET for researchers who did not have a computer on which to run CSNET software. The public host consists of a Digital Equipment Corporation VAX 11/750 running Berkeley's VAX UNIX operating system with user privileges limited to sending and receiving memos.

Headed by Farber (Delaware) and Hearn (the Rand Corporation), the PhoneNet group had established two mail relays serving over 24 PhoneNet sites as shown on the map. Ultimately, the location and number of relay sites will be determined by the number of sites that remain on PhoneNet.

The PhoneNet software has been transported from a Version 6 UNIX system to a Berkeley VAX UNIX system, and packaged for distribution. Work is underway to add a uucp connection, and to integrate the MMDF mail transport system with Berkeley's new sendmail transport system. (Currently, MMDF and sendmail both perform the same general functions.)

2.2 Phase Two—Telenet Connectivity and Nameserver Services

By fall 1982, Phase Two was well underway. Headed by Comer^{*} (Purdue), the protocol group designed and implemented an X.25 protocol interface for the TCP/IP protocol software that DARPA provided to CSNET [5]. It will form the backbone of the CSNET software for Telenet users. The implementation runs on a Digital Equipment VAX under Berkeley VAX UNIX, using an Interactive Systems Corporation INCard/X.25 device to connect the VAX to Telenet.

In a demonstration on June 15, 1982, a computer at the University of Wisconsin at Madison was connected to a computer at Purdue University over CTC Telenet. The protocol software, transparent to user-level software like the mailer, establishes and disconnects host-to-host X.25 connections over Telenet automatically when the TCP/IP layer needs to send datagrams. The interface software, under test by contractor

^{*}John T. Korb assumed Comer's duties beginning August 1, 1982 during Comer's sabbatical leave.

sites until late in 1982, becomes available for CSNET member sites in the fall of 1983.

Work has concentrated on finding cost-effective ways of using the X.25 service provided by Telenet with TCP/IP. A difference in basic philosophy between IP and X.25 accounts for the problem: IP uses datagrams while X.25 uses virtual circuits. With the datagram approach, the system divides a message into packets, places the destination address on each packet and sends the packet to the network for delivery, analogous to the way one sends telegrams. As the datagram moves across the network, its path is computed dynamically, depending on traffic. Virtual circuit services establish and record a path through the network to the destination machine (at a relatively high cost), and then send packets back and forth across the path (at a lower cost). When the user closes the virtual circuit, the network eliminates the record of the path. One cannot afford to open an X.25 virtual circuit for each datagram that must be sent because the charges for opening circuits are high. The IP-to-X.25 interface software manages a set of X.25 virtual circuits, automatically disconnecting a virtual circuit when it has been idle. The interface also multiplexes all communications between two machines over the same X.25 connection to minimize cost. (More details can be found in [5].)

Headed by Solomon (Wisconsin), the CSNET nameserver group has produced the Phase Two service software [6]. It runs on a second VAX 750, called the service host, located at the University of Wisconsin.

The nameserver look-up scheme uses a list of key words that the user supplies in addition to the user's name to help narrow the search. Authority to register a site rests with the CSNET Coordination and Information Center (CIC). Once a site entry has been established, the software allows users to install and edit their own entries. Users might choose to list their area of research, personal attributes, or even common misspellings of their names as keywords—any string that would help identify them. To send mail, one will only need to list enough keywords to identify the recipient.

By November 1982, a preliminary version of the nameserver running on the service host responded to interactive queries from CSNET machines on ARPANET or Telenet. Users can edit their own entries, and find the network address of colleagues interactively. Eventually, all CSNET mail systems, including the telephone-based relay systems, will consult the nameserver to determine mail addresses automatically, based on names and keywords in the header. When installed on member's systems, the automatic look-up facility will eliminate many delays caused by incorrect or unknown addresses.

2.3 Phase Three—Organization and Management

Because the Science Board restricted direct management by NSF to two years, much energy has been focused on forming organizations to take over CSNET. Two organizations were envisioned in the original plan: a Coordination and Information Center (CIC), which would manage the network; and a controlling organization, sometimes called CSNET Inc., which would have ultimate authority.

2.4 The Coordination and Information Center

According to the original plan, CIC would continue CSNET development after the initial two-year period; help new users establish connections; provide information to the public and to member sites; answer questions from a "hotline"; and collect and disperse monies from membership dues, equipment and software, and network use. CIC, responsible for day-to-

day planning, would also be able to investigate bulk rate discounts for network use. It would keep abreast of technology, and suggest improvements and enhancements to the controlling organization.

In the fall of 1981, NSF initiated its search for an organization to assume the role of CIC by issuing a request for proposals open to all qualified institutions (except the four CSNET contractors). The request described the responsibilities of the CIC, but permitted the proposing institution flexibility in choosing how it would fulfill them. From the review process, NSF selected Bolt Beranek and Newman (BBN) as the home of the CIC. At BBN, Richard Edmiston heads the CIC and sits on the management committee.

2.5 The Controlling Organization

As envisioned, the controlling organization, probably composed of a representative from each member site, would review progress, set policy, and give long range direction to the network. During the spring and summer of 1982, the organization and policy support groups together drew up a proposed constitution and a set of by-laws for the controlling organization. One of its responsibilities was using this material in preparing the solicitation. The results of the search will be announced by October 1983.

3. COSTS OF CSNET SERVICES

Potential CSNET sites commonly ask: "How much will a connection cost?" Of course, the answer depends on the type of connection, frequency and volume of traffic, and type of organization. NSF has approved a dues structure as shown in Table 1.

Estimates for hardware, connection, and traffic have been computed for various size organizations, based on traffic observed on ARPANET and traffic between early CSNET sites. (These are summarized in [7].) Telenet sites can expect annual costs (excluding dues) of between \$15,000 to \$21,000 for use by approximately 20 to 50 researchers. In addition, equipment costing approximately \$10,000 is needed to connect to Telenet. For sites with little traffic, PhoneNet is much less expensive: the annual charge for 20 researchers is estimated to be \$9,000. At higher traffic volumes, however, telephone connections become more expensive than Telenet connections. With 50 users, PhoneNet is expected to cost \$25,000.

4. SUMMARY AND CONCLUSION

4.1 Summary

CSNET is a logical network spanning several physical nets: ARPANET, Telenet, and a lowcost telephone-based mail relay service called PhoneNet. The goal is to provide services and connectivity among all persons engaged in computer science research. The cost of CSNET service will be proportional to the type and volume of use. Seed money, provided by NSF, will cover development costs but the network must be self-sufficient after five years. Development is carried out by groups at the University of Delaware, Purdue University, the

TABLE 1: Annual Dues for CSNET

Industrial site	\$30,000
Government or Non-Profit site	\$10,000
University site	\$5,000

Information on CSNET Membership

CSNET is open to institutions that support a significant computer research activity. Information can be obtained by writing:

CSNET CIC,
Bolt, Beranek, and Newman Inc.
10 Moulton Street
Cambridge, MA 02238

or by calling the CSNET Hotline (617) 497-2777.

Rand Corporation, and the University of Wisconsin, under NSF management. CIC, responsible for managing the network and distributing software, is located at Bolt, Beranek, and Newman Inc. A parent organization has not yet been chartered. A large portion of the computer science research community has been involved in the formation, planning, reviewing, and support of CSNET.

4.2 Conclusions

Perhaps it is too early to draw conclusions about CSNET. Some observers wonder whether CSNET can be self-sufficient and how it will keep up with changes in technology. Others talk of starting online technical journals, distributed research projects, and automatic software distribution over the network as if everyone already had access.

A larger question arises about the community's awareness of networks in general. It is obvious from the reviews of the CSNET proposal that many researchers remain unenlightened. Some said CSNET could not run over a public packet-switched network. (Now, it does.) Others said it had already been done. (It hadn't.) One reviewer even asserted that networks were unimportant because they are only used to make luncheon appointments. Vendors promised that their off-the-shelf systems would solve all the technical problems, and many reviewers believed it. (The earliest commercial network with similar services is years away).

Many sites are still naive about the cost of network connections as well. An ARPANET node (IMP) costs roughly \$107,000 annually in 1983 (flat rate, no additional traffic charges). By comparison, the cost estimate for a CSNET Telnet connection that supports 30 heavy users and 20 moderate users is \$21,000 annually. Thus, the benefits of network communications do not come for free.

My personal view of CSNET can be explained best by example. When I first met Kerri, he was preparing to present CSNET to the Science Board for NSF. He asked me whether I thought CSNET was worthwhile. "Yes," I told him, "connecting all computer scientists will change the way they do research." From the look on his face, I could tell he was not quite sure he agreed; computer science and computer networks were new and foreign to him. Recently, after he had used network facilities for over a year, I asked him whether he thought a network like CSNET would help chemists. He lit up. "Yes," he said, "I'm sold on network technology."

Acknowledgment. I would like to thank the many people who took time to answer questions and provide information; I regret that I could not mention every individual's contribution to the project. I would also like to thank the several reviewers who provided comments on an earlier draft. Finally, I thank CIC for preparing the map of CSNET sites.

REFERENCES

1. Comer, D.E. and Korb, J.T. CSNET protocol software: The IP-to-X.25 interface. *Proc. Symp. Data Comm. ACM SIGCOMM* (March 1983), 154-159.
2. Crocker, D.H., Szwarcowski, E.S., and Farber, D.J. An Inter-Network Memo Distribution Capability—MMDF. *Proc. Sixth Data Comm. Symp. ACM SIGCOMM* (1979), 18-25.
3. Denning, P.J., Hearn, A., and Kern, C.W. History and overview of CSNET. *Proc. Symp. Data Comm. ACM SIGCOMM* (March 1983), 138-145.
4. Feldman, J. and Sutherland, W. Rejuvenating experimental computer science—A report to the national science foundation and others. *Comm. ACM* 22, 9 (September 1979), 497-502.
5. From the meeting of the National Science Foundation Computer Science and Engineering Advisory Panel, December 5&6, 1974, cited in *NSF Program Report vol 5:6* (Oct. 1981), 18.
6. Landweber, L., Litzkow, L., Neuhengen, D., and Solomon, M. Architecture of the CSNET name server. *Proc. Symp. Data Comm. ACM SIGCOMM* (March 1983), 146-153.
7. Landweber, L. and Solomon, M. Use of Multiple Networks in CSNET. *Proc. IEEE COMPON* (February 1982), 146-153.

CR Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design; K.2 [Computing Milieux]: History of Computing

General Terms:

Additional Key Words and Phrases: CSNET, TCP/IP

Received 5/83; revised 7/83; accepted 8/83

COMING IN NOVEMBER ...

A Special Issue On
Research Contributions